Generalization of Meta Learning and Personalized Federated Learning
Master Thesis Project at LAS Lab, ETH Zurich

Yifan Hu*

March 1, 2023

This project aims to investigate the appropriate formulation for meta learning and personalized federated learning with a focus on their generalization performance. Specifically, we will explore the mathematical models that need to be solved, so that the obtained parameter will have good generalization performance on a new unseen task given limited samples.

Key words: generalization bounds, meta-learning, personalized federated learning, stability, sample complexity, gradient-based methods.

One sentence summary: We aim to find the correct formulation for many important problems in meta-learning and design algorithms to solve the correct formulation.

1 Introduction

Since 2017, optimization-based meta-learning has achieved great success in various fields of machine learning, including reinforcement learning, federated learning, computational biology, domain adaptation, and many others, with Finn et al. [2017] acquiring more than 8000 citations. The core idea of meta-learning is to learn a common shared parameter for several similar tasks. For instance, it can be a common shared initialization neural network parameter, from which a good parameter for a specific task can be obtained with just a few steps of gradient descent updates (see MAML [Finn et al., 2017]). It can also be common shared neural network parameters of the first $d-1$ layers, while the parameter of the last layer is trained using the data of a specific task (see representation learning [Bengio et al., 2013]).

Given the empirical success of meta-learning, researchers have started to study its theoretical perspectives, including generalization performance and optimization performance. Generalization performance refers to how the learned common knowledge can help a new unseen task with few data points to generalize. Optimization performance refers to how to design efficient optimization algorithms for finding the common shared parameters. By studying generalization and optimization, we can determine the correct mathematical models to solve and how fast we can obtain a good result. In this project, our goal is to study the generalization performance of several widely considered meta-learning formulations and determine which one provides the best generalization bounds for a new unseen task.

However, one key drawback of existing research is that the generalization bounds focus on how the learned common shared parameter generalizes on a meta-learning objective used for training purposes. In reality, what people care about is how a new unseen task can benefit from the learned common shared parameter and generalize well even with few data points. Therefore, our project aims to address this gap in the literature.

To gain insight into similar studies, we recommend referring to Fallah et al. [2021] to get an idea of our proposed project.

*LAS Lab at ETH and College of Management of Technology at EPFL, Switzerland. Emails: yifan.hu@inf.ethz.ch,

2 Logistics

The thesis project will primarily focus on theoretical analysis, and the selected student will work in close collaboration with Yifan Hu. The project timeline is as follows:

- One month dedicated to conducting a literature review on generalization bounds for meta-learning and formulating the problem.
- Two to three months devoted to theoretical analysis of generalization bounds for both meta-learning and personalized federated learning.
- One to two months allotted for designing algorithms for the proposed formulation and conducting experiments.
- One month dedicated to wrapping up the project and writing the master's thesis.

Upon completion of the master's thesis, the student will have the opportunity to present their findings at the group meeting of LAS.

Applicants with a strong mathematical and/or statistical background and experience coding with Python for deep learning using either TensorFlow or PyTorch are encouraged to apply. Interested individuals should contact Yifan directly at yifan.hu@inf.ethz.ch.

References

Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Generalization of model-agnostic meta-learning algorithms: Recurring and unseen tasks. *Advances in Neural Information Processing Systems*, 34:5469–5480, 2021.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017.